

Combining low-dimensional ensemble postprocessing with reordering methods

Roman Schefzik

*Heidelberg Institute for Theoretical Studies
Schloss-Wolfsbrunnengasse 35, 69118 Heidelberg, Germany*

roman.schefzik@h-its.org

Abstract

State-of-the-art weather forecasts usually rely on ensemble prediction systems, accounting for the different sources of uncertainty. As ensembles are typically uncalibrated, they should get statistically postprocessed. Several multivariate ensemble postprocessing techniques, which additionally consider spatial, inter-variable and/or temporal dependencies, have been developed. These can be roughly divided into two groups. The first group comprises parametric, mainly low-dimensional approaches that are tailored to specific settings. The second group involves non-parametric reordering methods that impose a specific dependence template on univariately postprocessed forecasts and are suitable in any dimension. In this paper, these different strategies are combined, with the aim to exploit the benefits of both concepts. Specifically, a high-dimensional postprocessing problem is divided into multiple low-dimensional instances, each of which is postprocessed via a suitable multivariate parametric method. From each postprocessed low-dimensional distribution, a sample is drawn, which is then reordered according to the corresponding multidimensional rank structure of an appropriately chosen dependence template. In this context, different ranking concepts for multivariate settings are discussed. Finally, all reordered samples are aggregated to obtain the overall postprocessed ensemble. The new approach is applied to ensemble forecasts for temperature and wind speed at several locations from the European Centre for Medium-Range Weather Forecasts, using a recent bivariate ensemble model output statistics postprocessing technique and a reordering based on the raw ensemble forecasts similar to the ensemble copula coupling method. It shows good predictive skill and outperforms reference ensembles.

Keywords: ensemble copula coupling, ensemble model output statistics, multivariate ranking, probabilistic weather forecasting, statistical ensemble postprocessing

1 Introduction

Nowadays, weather forecasts are typically derived from ensemble prediction systems, which comprise different runs of numerical weather prediction models varying in terms of the initial conditions and/or the parameterization of the atmosphere (Gneiting and Raftery, 2005; Leutbecher and Palmer, 2008; Palmer, 2002). While ensembles thus account for the main sources of predictive uncertainty, they tend to exhibit biases and lack calibration (Hamill and Colucci, 1997), in that the verifying observation

falls too often outside the ensemble range (Gneiting et al., 2007). Consequently, an ensemble forecast should get statistically postprocessed. Numerous ensemble postprocessing methods have been introduced during the last decade, with the ensemble model output statistics (EMOS) (Gneiting et al., 2005; Thorarinsdottir and Gneiting, 2010, among others) approach and Bayesian model averaging (BMA) (Baran, 2014; Raftery et al., 2005, among others) being two key examples. However, most of these techniques, including standard EMOS and BMA, do not account for potentially crucial dependence structures in space or time, or among different weather quantities.

To address this, a lot of work has been done on multivariate postprocessing approaches modeling correlation patterns adequately, in particular in recent years. Among others, variants of EMOS and BMA, respectively, which are able to deal with spatial (Berrocal et al., 2007, 2008; Feldmann et al., 2015) or inter-variable (Baran and Möller, 2015a,b; Möller et al., 2013; Schuhen et al., 2012; Sloughter et al., 2013) dependencies have been proposed. Moreover, temporal dependencies of consecutive lead times in postprocessed predictive distributions can be treated via the approaches of Pinson et al. (2009) or Schoelzel and Hense (2011). The aforementioned multivariate postprocessing techniques are parametric and perform well in low-dimensional settings, or when the involved correlation matrix can be assumed to be highly structured.

However, they mostly address either spatial or inter-variable or temporal dependencies only and are not suitable for high-dimensional settings without any specific structure. In such cases, multivariate ensemble postprocessing can be appropriately done using non-parametric reordering methods (Scheffzik, 2015b; Wilks, 2015), which are conceptionally simpler than the parametric approaches. Essentially, these methods rearrange for each margin a sample drawn from the respective predictive distribution obtained by univariate ensemble postprocessing (via EMOS or BMA, for instance) according to the rank dependence structure of a suitable dependence template (Scheffzik, 2015b; Wilks, 2015). Aggregating these reordered samples yields the postprocessed ensemble respecting dependence patterns. Examples of reordering methods include the Schaake shuffle (Clark et al., 2004), as well as specific implementations thereof (Scheffzik, 2015b), and ensemble copula coupling (ECC) (Scheffzik et al., 2013). While the Schaake shuffle employs a database of historical observations as a dependence template, ECC relies on the information in the raw ensemble forecast instead. From a mathematical point of view, reordering methods can be interpreted in the general frame of empirical copulas, as discussed by Scheffzik (2015a). In this paper, we develop a new approach which combines parametric Low-Dimensional Postprocessing (LDP) methods tailored to specific settings on the one hand with the non-parametric Reordering notion on the other hand, with the goal of exploiting the benefits of both concepts. It will be referred to as the LDP-Reordering approach in what follows. In the context of this new technique, the question of how to define a ranking concept in multivariate settings comes into play. The course of action in the paper is two-part. First, the LDP-Reordering method is presented in a very general frame, being broadly applicable. Second, we consider an example, in which the bivariate EMOS model of Baran and Möller (2015b) is used for the joint postprocessing of 10 meter wind speed and surface temperature forecasts at distinct stations individually. The respective postprocessed bivariate samples are then reordered and aggregated via an ECC-like approach using an appropriate reordering notion for multivariate settings. We refer to this specific implementation, which will be tested and evaluated in a case study exemplarily for an LDP-Reordering method, as the Bivariate EMOS-ECC approach.

The remainder of this paper is organized as follows. In view of the aforementioned illustrative example and the later case study, we review in Section 2 relevant univariate and bivariate EMOS approaches,

as well as the concept of reordering methods, with an emphasis on ECC. Moreover, notions of ranking in multivariate settings are discussed. In Section 3, we first introduce the general LDP-Reordering approach and then consider Bivariate EMOS-ECC as an example. Section 4 deals with the evaluation of the Bivariate EMOS-ECC approach in a case study using real forecast data from the European Centre for Medium-Range Weather Forecasts. Finally, the paper closes with a discussion in Section 5.

2 Preliminaries

In this preparatory section, we review existing postprocessing methods that are relevant in our case study. In addition, different ranking concepts in multivariate settings are elucidated.

2.1 Reference ensemble postprocessing methods

2.1.1 Ensemble model output statistics (EMOS)

The standard univariate ensemble model output statistics (EMOS) approach, also known as non-homogeneous regression, fits a single probability distribution for a fixed weather quantity at a fixed location and for a fixed look-ahead time, using summary statistics from a raw ensemble forecast x_1, \dots, x_M . The choice of the specific type of the predictive distribution depends on the considered weather variable. For instance, EMOS variants are available for temperature and pressure (Gneiting et al., 2005), wind speed (Baran and Lerch, 2015; Lerch and Thorarinsdottir, 2013; Thorarinsdottir and Gneiting, 2010) or precipitation (Scheuerer, 2014).

In the case of temperature y , EMOS employs a normal predictive distribution

$$\mathcal{N}(a + b_1x_1 + \dots + b_Mx_M, c + ds^2) \quad (1)$$

having mean $a + b_1x_1 + \dots + b_Mx_M$ and variance $c + ds^2$, where a, b_1, \dots, b_M, c and d are parameters that need to be estimated, and $s^2 := (1/M) \sum_{m=1}^M (x_m - \bar{x})^2$ denotes the empirical ensemble variance, with the empirical ensemble mean $\bar{x} := (1/M) \sum_{m=1}^M x_m$ (Gneiting et al., 2005). While a is a bias correction term, the regression coefficients b_1, \dots, b_M reflect the performance of the ensemble members over the training period relative to the other members, as well as the correlations between the ensemble members (Gneiting et al., 2005). If the ensemble members are statistically indistinguishable and can thus be considered exchangeable, the regression coefficients are assumed to be equal, such that $b_1 = \dots = b_M$. For wind speed y , which is a nonnegative quantity, we follow Thorarinsdottir and Gneiting (2010) and use a truncated normal distribution with a cut-off at zero, that is,

$$\mathcal{N}^0(a + b_1x_1 + \dots + b_Mx_M, c + ds^2), \quad (2)$$

with parameters as in (1).

Both in (1) and (2), the parameters a, b_1, \dots, b_M, c and d are estimated by optimizing a proper scoring rule (Gneiting and Raftery, 2007) over a training data set consisting of past forecasts and observations. According to Gneiting et al. (2005), common choices for the proper scoring rule to be optimized are the logarithmic (or ignorance) score (Good, 1952; Roulston and Smith, 2002), following the classical maximum likelihood notion, and the continuous ranked probability score (Gneiting and Raftery, 2007; Hersbach, 2000). For both scores, closed-form expressions are available in the above two cases.

The EMOS approach for temperature is implemented in the R package `ensembleMOS` (R Core Team, 2014; Yuen et al., 2013).

To get an ensemble representation of a full EMOS predictive cumulative distribution function (CDF) F , we draw a sample $\tilde{x}_1, \dots, \tilde{x}_N$ thereof. This can be for instance done by taking the equally spaced quantiles

$$\tilde{x}_1 := F^{-1}\left(\frac{1}{N+1}\right), \dots, \tilde{x}_N := F^{-1}\left(\frac{N}{N+1}\right). \quad (\text{Q})$$

Alternatively, we can just sample randomly from F , that is,

$$\tilde{x}_1 := F^{-1}(u_1), \dots, \tilde{x}_N := F^{-1}(u_N). \quad (\text{R})$$

where u_1, \dots, u_N are independent standard uniform random variates.

The corresponding ensembles are referred to as EMOS-Q and EMOS-R ensembles, respectively, in this paper, following a similar notation as in Schefzik et al. (2013).

2.1.2 Bivariate EMOS for joint postprocessing of wind speed and temperature forecasts

For the joint postprocessing of non-negative wind speed and real-valued temperature forecasts at a fixed location and for a fixed prediction horizon, Baran and Möller (2015b) propose to use a bivariate normal distribution $\mathcal{N}_2^0(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with the first coordinate truncated from below at zero, where $\boldsymbol{\mu}$ is a location vector and $\boldsymbol{\Sigma}$ a scale matrix. Precisely, if $\mathbf{x}_1 := (x_1^W, x_1^T), \dots, \mathbf{x}_M := (x_M^W, x_M^T)$ denotes the bivariate M -member raw ensemble forecast for wind speed (W) and temperature (T), the corresponding EMOS predictive distribution is given by

$$\mathcal{N}_2^0(\mathbf{A} + \mathbf{B}_1\mathbf{x}_1 + \dots + \mathbf{B}_M\mathbf{x}_M, \mathbf{C} + \mathbf{DSD}'), \quad (3)$$

where

$$\mathbf{S} := \frac{1}{M-1} \sum_{m=1}^M (\mathbf{x}_m - \bar{\mathbf{x}})(\mathbf{x}_m - \bar{\mathbf{x}})',$$

with $\bar{\mathbf{x}}$ being the empirical ensemble mean vector and \mathbf{Y}' denoting the transpose of the matrix or vector \mathbf{Y} . In (3), \mathbf{A} is a bivariate parameter vector, and $\mathbf{B}_1, \dots, \mathbf{B}_M, \mathbf{C}$ and \mathbf{D} are two-by-two real parameter matrices, with \mathbf{C} being symmetric and non-negative definite. Baran and Möller (2015b) estimate the parameters by optimizing with respect to the mean logarithmic score, using past ensemble forecasts and observations as training data. For details, see Baran and Möller (2015b).

To sample from a bivariate normal distribution with the first coordinate truncated from below at zero, one can use a rejection sampling method, which draws from an untruncated bivariate normal distribution and accepts only those samples that lie inside the support region. Alternatively, a Gibbs sampler can be employed. Both sampling options are implemented in the R package `tmvtnorm` (Wilhelm and Manjunath, 2010, 2014).

2.1.3 Reordering methods, in particular ensemble copula coupling (ECC)

The bivariate EMOS method from 2.1.2 joins the steadily growing list of purely parametric ensemble postprocessing techniques tailored to specific low-dimensional settings (Baran and Möller, 2015a;

Möller et al., 2013; Pinson, 2012; Schuhen et al., 2012, among others). However, parametric approaches are often inadequate in high-dimensional settings and fail to address spatial, inter-variable and temporal dependence structures simultaneously. In contrast, non-parametric methods based on reordering notions provide suitable and appealing tools to model such dependencies together, also in higher dimensions.

For the following general description of reordering methods, we stick to Scheme 1 in Schefzik (2015b). Let ℓ be a fixed margin, summarizing a fixed weather quantity, a fixed location and a fixed prediction horizon, and let L be the total number of margins. The goal is to get a postprocessed ensemble forecast $\hat{\mathbf{x}} := \{(\hat{x}_1^1, \dots, \hat{x}_N^1), \dots, (\hat{x}_1^L, \dots, \hat{x}_N^L)\}$ out of an M -member raw ensemble forecast $\mathbf{x} := \{(x_1^1, \dots, x_M^1), \dots, (x_1^L, \dots, x_M^L)\}$. In a first step, we choose a suitable data set $\mathbf{z} := \{(z_1^1, \dots, z_N^1), \dots, (z_1^L, \dots, z_N^L)\}$ which serves as a dependence template (Wilks, 2015) and adequately models correlation structures among the L margins. For each margin $\ell \in \{1, \dots, L\}$, we derive the univariate order statistics

$$z_{(1)}^\ell \leq \dots \leq z_{(N)}^\ell, \quad (4)$$

inducing the permutation $\pi_\ell(n) := \text{rank}(z_n^\ell)$ for $n \in \{1, \dots, N\}$, with ties resolved at random. Then, univariate postprocessing of the raw ensemble forecast $x_1^\ell, \dots, x_M^\ell$ is performed for each margin ℓ , for instance via EMOS, yielding a postprocessed predictive CDF F_ℓ . From each marginal CDF F_ℓ , we draw a sample $\tilde{x}_1^\ell, \dots, \tilde{x}_N^\ell$, for example by using one of the schemes (Q) or (R) from 2.1.1. Finally, we arrange these samples with respect to the rank dependence structure of \mathbf{z} . Hence, using the permutation π_ℓ , the final postprocessed ensemble $\hat{x}_1^\ell, \dots, \hat{x}_N^\ell$ for each margin ℓ is given by

$$\hat{x}_1^\ell := \tilde{x}_{(\pi_\ell(1))}^\ell, \dots, \hat{x}_N^\ell := \tilde{x}_{(\pi_\ell(N))}^\ell.$$

Aggregating the L postprocessed margins yields the desired overall postprocessed ensemble $\hat{\mathbf{x}}$. From a mathematical point of view, reordering methods can be interpreted in the frame of discrete or empirical copulas (Schefzik, 2015a,b; Wilks, 2015).

There are several options for the specific choice of the dependence template \mathbf{z} .

In the Schaake shuffle (Clark et al., 2004), \mathbf{z} is based on past observations from a historical database. A variant of the Schaake shuffle particularly employs historical observations from dates on which the corresponding ensemble forecast resembled the current one (Schefzik, 2015b).

In the ECC approach (Schefzik et al., 2013), \mathbf{z} is given by the raw ensemble forecast, that is, $\mathbf{z} = \mathbf{x}$. Hence, the size of the final postprocessed ensemble equals that of the raw ensemble, that is, $N = M$. ECC implicitly assumes that the ensemble is able to represent actual spatial, inter-variable and temporal dependence structures appropriately. This is surely not valid each and every day, but may largely be expected. Furthermore, ECC only applies to ensembles whose members are statistically indistinguishable and thus can be considered exchangeable.

In the special case of using univariate postprocessing via EMOS, together with the discretization schemes (Q) and (R), respectively, in the ECC approach, we refer to the corresponding ensembles as EMOS-ECC-Q and EMOS-ECC-R.

2.2 Ranking in multivariate settings

Ranking N one-dimensional data points $z_1, \dots, z_N \in \mathbb{R}$ can be conveniently done by using the usual univariate ordering $z_{(1)} \leq \dots \leq z_{(N)}$ as in (4), defining the permutation $\pi(n) := \text{rank}(z_n)$ for $n \in$

$\{1, \dots, N\}$, with ties resolved at random. The situation in the multivariate case is however more involved.

To rank N data points $\mathbf{z}_1 := (z_1^1, \dots, z_1^L), \dots, \mathbf{z}_N := (z_N^1, \dots, z_N^L) \in \mathbb{R}^L$ with L -dimensional output each, we proceed in the following steps:

Scheme 1. (Ranking in multivariate settings)

1. Each data point \mathbf{z}_n with index $n \in \{1, \dots, N\}$ is assigned a suitable multivariate characteristic $R_n \in \mathbb{R}$. There are several options to choose R_n , as will be discussed below.
2. We rank the values R_1, \dots, R_N by using the common univariate ordering $R_{(1)} \leq \dots \leq R_{(N)}$, inducing the permutation $\tau_n := \text{rank}(R_n)$ for $n \in \{1, \dots, N\}$, with ties resolved at random.
3. Each data point \mathbf{z}_n for $n \in \{1, \dots, N\}$ is finally assigned the corresponding rank τ_n derived in step 2.

As mentioned, there are several possibilities to define the multivariate characteristic R_n the ranking is based on. For instance, we can use the multivariate pre-rank construction following Gneiting et al. (2008) and set

$$R_n^{\text{mult}} := \sum_{\nu=1}^N \mathbb{1}_{\{\mathbf{z}_\nu \preceq \mathbf{z}_n\}}, \quad (5)$$

where $\mathbf{z}_\nu \preceq \mathbf{z}_n$ if and only if $z_\nu^\ell \leq z_n^\ell$ for all $\ell \in \{1, \dots, L\}$, and $\mathbb{1}_A$ denotes the indicator function of the event A which is 1 if A occurs and 0 otherwise. This concept works well in low-dimensional settings, as in the bivariate scenario in our case study, but may face problems when dealing with high dimensions (Thorarinsdottir et al., 2015). Another option to define a multivariate characteristic may employ the average pre-rank as introduced by Thorarinsdottir et al. (2015), which is especially suitable in high-dimensional settings. The average pre-rank

$$R_n^{\text{avg}} := \frac{1}{L} \sum_{\ell=1}^L \text{rank}_S(z_n^\ell) \quad (6)$$

is just the average over the univariate ranks, with $\text{rank}_S(z_n^\ell) = \sum_{\nu=1}^N \mathbb{1}_{\{z_\nu^\ell \leq z_n^\ell\}}$ denoting the rank of the ℓ -th coordinate of \mathbf{z}_n in $S := \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ for $n \in \{1, \dots, N\}$ (Thorarinsdottir et al., 2015). Further, alternative possibilities to define a multivariate characteristic may be based on minimum spanning trees (Gombos et al., 2007; Smith and Hansen, 2004; Wilks, 2004) or the band depth pre-rank (Thorarinsdottir et al., 2015).

While the above suggestions are rather general, we propose in the following a multivariate characteristic that is especially tailored to the setting of non-negative wind speed and real-valued temperature data in our case study. For wind speed (W) data points $z_1^W, \dots, z_N^W \in \mathbb{R}_0^+$ and temperature (T) data points $z_1^T, \dots, z_N^T \in \mathbb{R}$, which should ideally all be standardized, we define what can be considered a signed Euclidean norm as follows:

$$R_n^{\text{Eucl}} := \text{sgn}(z_n^T) \cdot \sqrt{(z_n^W)^2 + (z_n^T)^2} \quad (7)$$

for $n \in \{1, \dots, N\}$, where

$$\text{sgn}(z_n^T) = \begin{cases} 1 & \text{if } z_n^T \geq 0 \\ -1 & \text{if } z_n^T < 0. \end{cases}$$




Figure 1: Illustration of multivariate ranking based on three different multivariate characteristics as discussed in Subsection 2.2, namely the multivariate pre-rank as in (5), the average pre-rank as in (6) and the signed Euclidean norm as in (7). Each data point indicated by a dot is labeled in black with its index and its corresponding multivariate characteristic in brackets. The number in gray indicates the final rank.

Again, this concept is specifically tailored to the setting of our case study and thus should not be used without further reflection as a general option in other settings.

An illustration of the multivariate ranking approaches based on the multivariate characteristics (5), (6) and (7), respectively, is given in Figure 1. In each plot, we show $N = 5$ data points $\mathbf{z}_1 := (z_1^1, z_1^2), \dots, \mathbf{z}_5 := (z_5^1, z_5^2)$ of dimension $L = 2$ each. With regard to the case study, we can think of $\mathbf{z}_1, \dots, \mathbf{z}_5$ as ensemble member forecasts, where the z_n^1 's can be considered temperature (T) forecasts, and the z_n^2 's wind speed (W) forecasts for $n \in \{1, \dots, 5\}$, thus $z_n^1 = z_n^T$ and $z_n^2 = z_n^W$. Each data point \mathbf{z}_n indicated by a dot is in each case labeled in black with its corresponding index n , ranging from 1 to 5, and its associated multivariate (bivariate) characteristic R_n in brackets. The final rank τ_n of each data point \mathbf{z}_n according to Scheme 1, ranging from 1 to 5 and obtained by ranking the multivariate characteristics according to common univariate ordering, is labeled in gray.

The plot at left in Figure 1 shows the ranking following the multivariate pre-rank idea in (5). For instance, the data point \mathbf{z}_5 with the index 5 has bivariate pre-rank $R_5^{\text{mult}} = 4$ according to (5), as 4 of the 5 data points lie to its lower left, including \mathbf{z}_5 itself. As this is the highest among all bivariate pre-ranks, \mathbf{z}_5 is assigned the final rank 5.

In the middle plot, the ranking based on the use of average pre-ranks according to (6) is illustrated. Exemplarily, the data point \mathbf{z}_5 has an average pre-rank of 4.5, as z_5^1 has rank 5 within the set $\{z_1^1, \dots, z_5^1\}$ and z_5^2 has rank 4 within the set $\{z_1^2, \dots, z_5^2\}$, such that $R_5^{\text{avg}} = (5 + 4)/2 = 4.5$. Again, this is the highest value among all pre-ranks, such that \mathbf{z}_5 is assigned rank 5.

Finally, the plot at right exemplifies the ranking procedure using the signed Euclidean norm in (7), in which the distance from the data point \mathbf{z}_n to the origin $O := (0, 0)$ is combined with a positive or negative sign, depending on whether $z_n^1 \geq 0$ or $z_n^1 < 0$, respectively. For example, the data point \mathbf{z}_5 has a signed Euclidean norm of 3.61, as the distance from \mathbf{z}_5 to O is 3.61 and $z_5^1 \geq 0$. In contrast, \mathbf{z}_1

has a signed Euclidean norm of -1.12 , since the distance from \mathbf{z}_1 to O is 1.12 and $z_1^1 < 0$. As \mathbf{z}_1 and \mathbf{z}_5 have the third highest and the highest signed Euclidean norms, respectively, they are assigned the final ranks 3 and 5, respectively.

Generally, for the signed Euclidean norm in (7), it is likely to obtain comparably few ties, if any, which reduces the need of a possible randomization in the ranking. In contrast, the chance of getting ties when using the multivariate pre-rank according to (5) appears to be higher, which also holds maybe to some lesser extent for the average pre-ranks in (6).

3 Combining Low-Dimensional Postprocessing (LDP) and Reordering methods

In this section, we present a new ensemble postprocessing approach which combines parametric Low-Dimensional Postprocessing (LDP) methods and the non-parametric Reordering notion from 2.1.3. The method applies a multivariate reordering to low-dimensional building blocks and will be referred to as the LDP-Reordering approach in what follows. The goal is to exploit the benefits of both concepts with respect to an adequate modeling of dependence structures. We first describe the LDP-Reordering method in a general frame being broadly applicable. In the second part of the section, we specifically discuss an example of the LDP-Reordering method dealing with the joint postprocessing of wind speed and temperature forecasts at different locations simultaneously. In this setting, which will also be used in the case study in Section 4, we combine the bivariate EMOS approach of Baran and Möller (2015b) from 2.1.2 with the ECC reordering notion of Schefzik et al. (2013) from 2.1.3 and refer to this procedure as the Bivariate EMOS-ECC method.

3.1 The LDP-Reordering method

We first introduce the general LDP-Reordering method. Let $i \in \{1, \dots, I\}$ be a weather variable, $j \in \{1, \dots, J\}$ a location and $k \in \{1, \dots, K\}$ a look-ahead time, summarized in the multi-index $\ell := (i, j, k)$, and let $L := I \times J \times K$. Moreover, let

$$\mathbf{x} := (\mathbf{x}_1, \dots, \mathbf{x}_M)$$

be the overall M -member raw ensemble forecast which is to be postprocessed, where each ensemble member $m \in \{1, \dots, M\}$ comprises an L -dimensional output

$$\mathbf{x}_m := (x_m^1, \dots, x_m^L).$$

In cases comprising weather variables of different magnitudes and/or units, one should work with standardized forecast values in what follows. Suppose that \mathbf{x} and \mathbf{x}_m for each ensemble member m , respectively, is split up into C sub-raw ensemble forecasts

$$\mathbf{x}_m^c := (x_m^{\ell^{c,1}}, \dots, x_m^{\ell^{c,L_c}})$$

where $c \in \{1, \dots, C\}$, such that

$$\mathbf{x}_m = (\mathbf{x}_m^1, \dots, \mathbf{x}_m^C).$$

In this context, $\ell^{c,1}, \dots, \ell^{c,L_c}$ denote the L_c multi-indices referring to those instances (weather variables, locations and/or look-ahead times) that shall be jointly postprocessed via a specific L_c -dimensional method in case $c \in \{1, \dots, C\}$. In addition, C is the number of cases in which a specific low-dimensional postprocessing method is intended to be applied, where the same approach may be used multiple times. Moreover, L_c is the dimension of the predictive CDF F_c obtained by the corresponding postprocessing method in the case $c \in \{1, \dots, C\}$. Thus, the overall dimension L can be decomposed via

$$L = \sum_{c=1}^C L_c.$$

The specific choice of the involved LDP methods should be as plausible as possible and tailored to the setting one deals with. It might be reasonable to postprocess variables with an expected pronounced correlation jointly. For example, at individual stations, wind vectors could be postprocessed simultaneously via the bivariate EMOS approach of Schuhen et al. (2012) or the method of Pinson (2012), while temperature and pressure could be treated jointly via the Gaussian copula approach of Möller et al. (2013). On the other hand, purely spatial settings could be handled by Spatial BMA (Berrocal et al., 2007) or Spatial EMOS (Feldmann et al., 2015), respectively. In our case study, we jointly postprocess wind speed and temperature forecasts via the bivariate EMOS method of Baran and Möller (2015b). This specific example has been described in detail in 2.1.2.

The general LDP-Reordering approach considered in this subsection proceeds in the following steps, with the goal to come up with a coherent, postprocessed ensemble forecast

$$\hat{\mathbf{x}} := (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N),$$

with

$$\hat{\mathbf{x}}_n := (\hat{\mathbf{x}}_n^1, \dots, \hat{\mathbf{x}}_n^C)$$

for each postprocessed member $n \in \{1, \dots, N\}$, where N in general does not need to be equal to the raw ensemble size M .

1. We choose a dependence template

$$\mathbf{z} := (\mathbf{z}_1, \dots, \mathbf{z}_N)$$

based on N data points, each of dimension L , where

$$\mathbf{z}_n := (z_n^1, \dots, z_n^L)$$

for an index $n \in \{1, \dots, N\}$. Similarly as above, we partition by setting

$$\mathbf{z}_n^c := (z_n^{\ell^{c,1}}, \dots, z_n^{\ell^{c,L_c}})$$

for case $c \in \{1, \dots, C\}$, such that

$$\mathbf{z}_n := (\mathbf{z}_n^1, \dots, \mathbf{z}_n^C),$$

where $n \in \{1, \dots, N\}$.

The following steps 2 to 5 are applied to each case $c \in \{1, \dots, C\}$ separately.

2. For each case c , we rank the L_c -dimensional points $\mathbf{z}_1^c, \dots, \mathbf{z}_N^c$ according to the scheme in Subsection 2.2. That is, we choose first a suitable multivariate characteristic R_n^c (with suggestions having been made in 2.2), where the corresponding choice has to be consistent for all involved cases. Then, we derive the permutation τ^c given by $\tau^c(n) := \text{rank}(R_n^c)$ for $n \in \{1, \dots, N\}$ using the usual univariate ordering, where ties are resolved at random.
3. We apply a reasonably chosen LDP method for the case c to obtain an L_c -variate postprocessed predictive CDF F_c .
4. From the predictive CDF F_c , we draw an L_c -dimensional sample $\tilde{\mathbf{x}}_1^c, \dots, \tilde{\mathbf{x}}_N^c$ of size N , where

$$\tilde{\mathbf{x}}_n^c := \left(\tilde{x}_n^{\ell^{c,1}}, \dots, \tilde{x}_n^{\ell^{c,L_c}} \right)$$

for $n \in \{1, \dots, N\}$.

5. We reorder the sample $\tilde{\mathbf{x}}_1^c, \dots, \tilde{\mathbf{x}}_N^c$ from the previous step by using the permutation τ^c determined in step 2 to get the ensemble $\hat{\mathbf{x}}_1^c, \dots, \hat{\mathbf{x}}_N^c$, with

$$\hat{\mathbf{x}}_n^c := \left(\hat{x}_n^{\ell^{c,1}}, \dots, \hat{x}_n^{\ell^{c,L_c}} \right)$$

for $n \in \{1, \dots, N\}$. To this end, we derive by analogy with step 2 the multivariate characteristics $\tilde{R}_1^c, \dots, \tilde{R}_N^c$ associated to $\tilde{\mathbf{x}}_1^c, \dots, \tilde{\mathbf{x}}_N^c$. This induces the corresponding order statistics $\tilde{R}_{(1)\sim}^c \leq \dots \leq \tilde{R}_{(N)\sim}^c$, with any ties resolved at random, using the notation $(\cdot)_{\sim}$ to emphasize that the ordering is based on the sample values. The corresponding permutation $\tilde{\tau}^c$ of $\{1, \dots, N\}$ is then given by $\tilde{\tau}^c(n) := \text{rank}(\tilde{R}_n^c)$, and the reordered ensemble $\hat{\mathbf{x}}_1^c, \dots, \hat{\mathbf{x}}_N^c$ is finally obtained via

$$\hat{\mathbf{x}}_n^c := \tilde{\mathbf{x}}_{(\tilde{\tau}^c)^{-1}(\tau^c(n))}^c$$

for $n \in \{1, \dots, N\}$. In other words using order statistic notation, $\hat{\mathbf{x}}_1^c, \dots, \hat{\mathbf{x}}_N^c$ is given by

$$\hat{\mathbf{x}}_n^c := \tilde{\mathbf{x}}_{(\tau^c(n))_{\sim}}^c$$

for $n \in \{1, \dots, N\}$. That is, the sample member with the index $(\tilde{\tau}^c)^{-1}(\tau^c(n))$ from step 4 forms the member n in the reordered ensemble. Therefore, the L_c -dimensional sample members $\tilde{\mathbf{x}}_1^c, \dots, \tilde{\mathbf{x}}_N^c$ from step 4 stay as they are with respect to their components, and only their order among themselves and hence the indices $n \in \{1, \dots, N\}$ are changed, according to the multivariate ranking structure of the chosen dependence template.

Having applied steps 2 to 5 to each case, we eventually aggregate all cases $c \in \{1, \dots, C\}$ to get the final LDP-Reordering ensemble $\hat{\mathbf{x}}$.

6. We first aggregate the postprocessed ensemble forecasts $\hat{\mathbf{x}}_n^c$ obtained in step 5 for all cases $c \in \{1, \dots, C\}$ for each member $n \in \{1, \dots, N\}$, which is done by setting

$$\hat{\mathbf{x}}_n := (\hat{\mathbf{x}}_n^1, \dots, \hat{\mathbf{x}}_n^C).$$

Finally, the LDP-Reordering ensemble $\hat{\mathbf{x}}$ is obtained via $\hat{\mathbf{x}} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N)$.

If a postprocessing method in case c is univariate, such as standard EMOS as in 2.1.1, the procedure reduces to a standard reordering approach.

For univariate methods, it may be most convenient and advisable to use equally spaced quantiles in the style of scheme (Q) in 2.1.1 in step 4. However, a generalization of the concept of equidistance to the multivariate case and its specific implementation are not straightforward. In our case study later, we thus employ random samples from the corresponding multivariate predictive distributions. The development of optimal sampling schemes for multivariate settings remains an important issue for future work.

For scenarios in which only one specific multivariate method is employed in a single case, the reordering in step 5 has no effect and thus becomes irrelevant. In these situations, our approach reduces to applying the corresponding LDP method and drawing a sample from the obtained multivariate predictive CDF. Finally, we recall that the established reordering methods for univariate margins as discussed in 2.1.3 can be mathematically interpreted in terms of empirical copulas (Scheffzik, 2015a). A similar copula-based framework for the LDP-Reordering method is very likely to hold, while explicit details remain to be worked out.

3.2 Specific example: Joint postprocessing of wind speed and temperature forecasts at several locations simultaneously using Bivariate EMOS and an ECC-like reordering

To give an example of an LDP-Reordering approach, we describe the scenario of the real-data case study using the $M = 50$ -member European Centre for Medium-Range Weather Forecasts (ECMWF) ensemble (Buizza, 2006; ECMWF Directorate, 2012; Molteni et al., 1996) whose results are presented in Section 4. We consider 24 hour ahead predictions for 10 meter wind speed (W) and surface temperature (T) at Vienna (Austria; Vie24), Bratislava (Slovakia; Br24) and Budapest (Hungary; Bu24), such that $I = 2$, $J = 3$ and $K = 1$, thus facing an $L = I \times J \times K = 6$ -dimensional setting. The bivariate EMOS approach of Baran and Möller (2015b) as described in 2.1.2 is applied to postprocess the wind speed and temperature forecasts jointly at the individual stations of Vienna, Bratislava and Budapest, respectively. Following the partitioning frame of the general LDP-Reordering method, the forecast \mathbf{x}_m of each raw ensemble member $m \in \{1, \dots, M\}$ is hence made up of

$$\begin{aligned} \mathbf{x}_m &:= (\mathbf{x}_m^{\text{Vie24}}, \mathbf{x}_m^{\text{Br24}}, \mathbf{x}_m^{\text{Bu24}}) \\ &:= (\underbrace{(x_m^{T, \text{Vie24}}, x_m^{W, \text{Vie24}})}_{\text{case 1 with } L_1=2}, \underbrace{(x_m^{T, \text{Br24}}, x_m^{W, \text{Br24}})}_{\text{case 2 with } L_2=2}, \underbrace{(x_m^{T, \text{Bu24}}, x_m^{W, \text{Bu24}})}_{\text{case 3 with } L_3=2}), \end{aligned}$$

thus including $C = 3$ sub-raw ensemble forecasts or cases, respectively, of dimension $L_c = 2$ each, where $c \in \{1, 2, 3\}$.

We later use the whole $M = 50$ -member ensemble in the evaluation of our method over a longer test period. However, for reasons of clarity and comprehensibility, we demonstrate the procedure for a subset of only $M = 8$ raw ensemble members by just taking the first 8 members of the 50-member ECMWF ensemble as the underlying raw ensemble in the following illustration. For the three locations, the left panels of Figure 2 show for each member $m \in \{1, \dots, 8\}$ the respective raw ensemble forecast values \mathbf{x}_m valid 1200 Universal Coordinated Time (UTC) on 14 December 2013, indicated

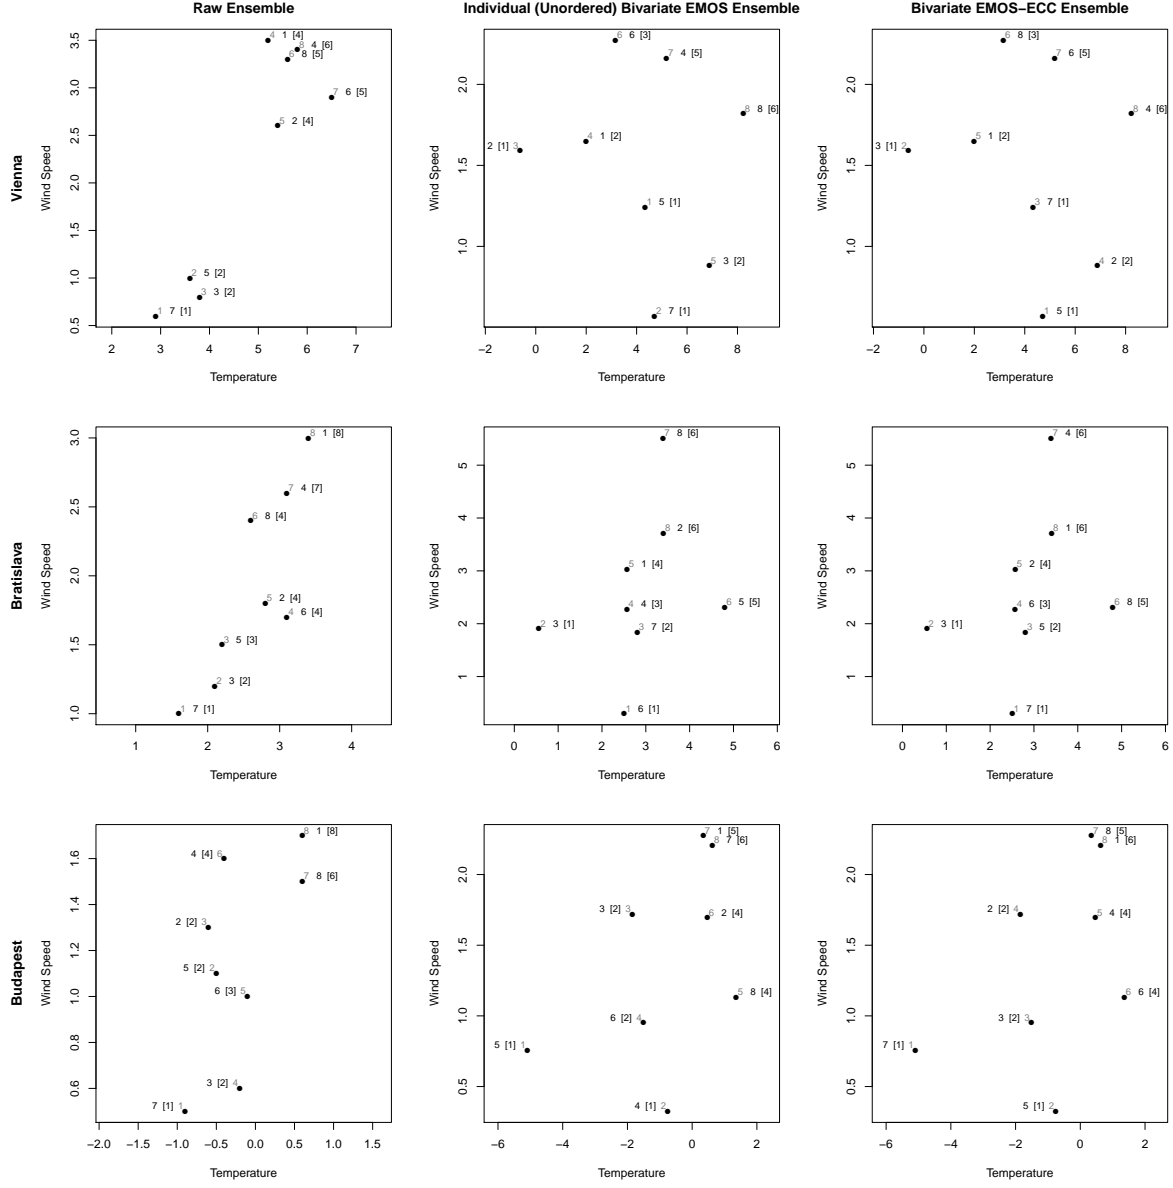


Figure 2: Illustration of the Bivariate EMOS-ECC approach based on the example in Subsection 3.2, valid 14 December 2013 , 1200 UTC. Each of the $M = 8$ ensemble members indicated by a dot is labeled in black with its index and its corresponding bivariate pre-rank according to (5) in brackets. The final rank of each member is given by the gray label. For reasons of clarity, the scales on the axes are different.

by the dots. Each member is labeled in black with its corresponding index and, in brackets, with its bivariate pre-rank according to (5), which is exemplarily used as a multivariate characteristic to define a bivariate ranking in the sense of Scheme 1 in Subsection 2.2. The final rank of each raw ensemble member is indicated by the corresponding label in gray.

The joint postprocessing of the wind speed and temperature ensemble forecasts at each location separately according to the bivariate EMOS method of Baran and Möller (2015b) yields bivariate predictive CDFs F_{Vie24} , F_{Br24} and F_{Bu24} , respectively. Using the rejection sampling algorithm implemented in the R package `tvmtvnorm` (Wilhelm and Manjunath, 2010, 2014), we generate a bivariate sample of size $M = 8$ from F_{Vie24} , F_{Br24} and F_{Bu24} , respectively. The mid-panel in Figure 2 shows these samples, which are referred to as the Individual (Unordered) Bivariate EMOS ensemble, again labeled in black with their corresponding indices and bivariate pre-ranks (in brackets) according to (5), which are derived analogously as for the raw ensemble, and labeled in gray with their final ranks.

The desired reordered postprocessed ensemble is finally based on an ECC-like notion using the raw ensemble as a dependence template, that is, we have $\mathbf{z} = \mathbf{x}$ and $N = M$ in the general LDP-Reordering approach introduced in Subsection 3.1. This ensemble is referred to as the Bivariate EMOS-ECC ensemble and is shown in the right panel of Figure 2. Again, each ensemble member is labeled in black with its index and its bivariate pre-rank (in brackets), and in gray with its final rank.

The pairs of ensemble member forecasts for wind speed and temperature forecasts and their corresponding bivariate pre-ranks (labeled in black in brackets), and thus also their respective final ranks (labeled in gray), for Vienna, Bratislava and Budapest, respectively, are exactly the same for both the Individual (Unordered) Bivariate EMOS (middle) and the Bivariate EMOS-ECC (right) ensemble. However, the members of the Bivariate EMOS-ECC approach have other indices than those of the Individual (Unordered) Bivariate EMOS ensemble, following the bivariate pre-rank structure in the raw ensemble. A final Bivariate EMOS-ECC ensemble member m for Vienna, Bratislava and Budapest, respectively, should have a bivariate pre-rank that is near to or comparable to the bivariate pre-rank which the raw ensemble member with the index m has. The same should consequently hold for the final rank then. An exact conformity of the bivariate pre-rank structure and the final rank pattern, respectively, of the raw and the Bivariate EMOS-ECC ensemble is not always possible, as in the samples drawn from the Bivariate EMOS distributions, other bivariate pre-ranks with different frequencies than in the raw ensemble might occur. However, the Bivariate EMOS-ECC approach tries to conserve the given structure to a great extent.

For instance, in our example for Budapest in the third row of Figure 2, the raw ensemble member with the index 7 has the lowest bivariate pre-rank (left panel), namely 1, and thus the final rank 1. This is respected by the Bivariate EMOS-ECC ensemble (right panel), whose member with the index 7 also has a bivariate pre-rank of 1 and hence also a final rank of 1. In contrast, this property gets drastically lost for the Individual (Unordered) Bivariate EMOS ensemble (mid-panel), for which the member with the index 7 has a bivariate pre-rank of 6, which is the highest in this case, and thus also the highest final rank, namely 8. So both the bivariate pre-rank and the final rank of the Individual (Unordered) Bivariate EMOS ensemble member with index 7 are unsuitably high with respect to the raw ensemble.

4 Case study

4.1 Setting and implementation details

We now evaluate the Bivariate EMOS-ECC approach in a case study, following the setting in the example in Subsection 3.2. That is, we consider 24 hour ahead joint 10 meter wind speed and surface temperature forecasts, all valid 1200 UTC, at Vienna (Austria), Bratislava (Slovakia) and Budapest (Hungary) simultaneously, provided by the $M = 50$ -member European Centre for Medium-Range Weather Forecasts (ECMWF) core ensemble (Buizza, 2006; ECMWF Directorate, 2012; Molteni et al., 1996). Our test period ranges from 1 January 2003 to 31 December 2013, comprising in total 3983 instances after the removal of dates with missing forecast and/or observation data. The temperature observations show very high spatial correlations among the three locations, while the spatial dependencies of the wind speed observations are highly to moderately pronounced. Inter-variable correlations between the temperature and wind speed observations are existent, but are not that marked compared to the spatial dependencies. All observational correlation structures are mirrored reasonably well in the raw ensemble.

We employ both univariate and bivariate EMOS methods to postprocess the wind speed and temperature predictions. In both cases, the respective model parameters are estimated based on training data only from the individual station of interest, thus yielding a distinct set of parameters for each location. We use a sliding window of 50 days as a training period for parameter estimation in both implementations.

Univariate postprocessing, which applies to each weather variable separately, is performed via the EMOS approaches (1) for temperature and (2) for wind speed, respectively, as discussed in 2.1.1. Optimization in the univariate EMOS approaches is performed with respect to the continuous ranked probability score (CRPS), because if a closed-form expression of the CRPS is available, which is the case for the models (1) and (2), this usually yields better results than optimizing with respect to the logarithmic score. Concerning the joint postprocessing of wind speed and temperature forecasts, we use the bivariate EMOS approach of Baran and Möller (2015b) as presented in 2.1.2. For this method, we follow Baran and Möller (2015b) and estimate the model parameters by optimizing the mean logarithmic score, as there is no closed-form expression available for the CRPS in the case of a bivariate truncated normal distribution.

As we aim to compare the different approaches in form of ensemble forecasts, we need to discretize the (either univariate or bivariate) postprocessed predictive distributions. In the univariate case, we employ the EMOS-R ensemble, which is generated by using quantization scheme (R) from 2.1.1, that is, we just draw a random sample from each marginal predictive distribution. Similarly, the samples from the Bivariate EMOS predictive distributions are also drawn in a random manner by using the rejection sampling method implemented in the R package `tvmvtnorm` (Wilhelm and Manjunath, 2010, 2014). Note that for univariate settings, samples according to scheme (Q) in 2.1.1 are expected to yield better results than samples based on scheme (R), as equidistant quantiles can be considered kind of an optimal choice for univariate sampling (Bröcker, 2012; Graf and Luschgy, 2000). In contrast, there appears to be no obvious optimal sampling strategy for multivariate settings, such that for a fairer comparison, we only consider random sampling both in the univariate and the bivariate case here.

Having performed the discretization, the (uni- or bivariate) samples get reordered according to the dependence structure in the raw ensemble. In the univariate case, this is done by applying the ECC

approach by Schefzik et al. (2013) as discussed in 2.1.3, finally leading to the EMOS-ECC-R ensemble. For the Bivariate EMOS samples, we perform the reordering via the ECC-like procedure illustrated in Subsection 3.2, where we study variants based on the different bivariate ranking characteristics according to the bivariate pre-rank (BivPR) in (5), the average pre-rank (AvPR) in (6) and the signed Euclidean norm (SEN) in (7) as described in Subsection 2.2. We refer to the corresponding ensembles as the Bivariate EMOS-ECC ensemble (BivPR), the Bivariate EMOS-ECC ensemble (AvPR) and the Bivariate EMOS-ECC ensemble (SEN), respectively, in what follows.

In a nutshell, we compare the predictive performances of the following reference ensembles:

- the unprocessed raw ensemble,
- the EMOS-ECC-R ensemble,
- the Individual (Unordered) Bivariate EMOS ensemble consisting of the unordered bivariate samples obtained from the Bivariate EMOS postprocessing, by analogy with the situation illustrated in the mid-panel of Figure 2, and
- three different Bivariate EMOS-ECC ensembles, namely
 - the Bivariate EMOS-ECC ensemble (BivPR),
 - the Bivariate EMOS-ECC ensemble (AvPR), and
 - the Bivariate EMOS-ECC ensemble (SEN).

4.2 Evaluation tools

Following Gneiting et al. (2007), a predictive distribution or an ensemble forecast, respectively, should be as sharp as possible, but subject to being calibrated, where calibration refers to the statistical consistency between forecasts and observation. There are several verification tools to evaluate the predictive performances of our reference ensembles (Wilks, 2011). To describe the evaluation tools used in this paper, let

$$\mathbf{x}_1 := (x_1^1, \dots, x_1^L), \dots, \mathbf{x}_N := (x_N^1, \dots, x_N^L) \in \mathbb{R}^L$$

denote an N -member ensemble forecast and

$$\mathbf{y} := (y_1, \dots, y_L) \in \mathbb{R}^L$$

an observation in the case of an individual forecast instance.

We assess calibration by computing the rank of the observation \mathbf{y} within the pooled set $\{\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{y}\}$ for all individual forecast cases and plotting the respective rank histogram. In the case of a calibrated ensemble forecast, the ranks are uniformly distributed on $\{1, \dots, N + 1\}$, thus yielding a flat rank histogram. In univariate settings, the common ranking can be employed, leading to the verification rank histogram (Anderson, 1996; Hamill, 2001; Talagrand et al., 1997). For multivariate quantities, which are our focus, we use the multivariate, band depth and average ranking concepts and histograms,

respectively (Gneiting et al., 2008; Thorarinsdottir et al., 2015). A quantitative measure for the deviation from a uniform distribution on $\{1, \dots, N+1\}$ is the reliability index Δ given by

$$\Delta := \sum_{r=1}^{N+1} \left| \rho_r - \frac{1}{N+1} \right|,$$

where ρ_r denotes the relative frequency of rank r (Delle Monache et al., 2006). Depending on if we use the multivariate, band depth or average rank concept for the ranking, we refer to Δ_{MR} , Δ_{BDR} and Δ_{AvR} , respectively, for the corresponding histograms in what follows.

The overall forecast skill is evaluated via proper scoring rules (Gneiting and Raftery, 2007), which give a single number assessing calibration and sharpness simultaneously. We take scoring rules to be negatively oriented here, meaning that the lower the score the better the predictive performance. An established proper scoring rule for univariate quantities is the continuous ranked probability score (CRPS) (Gneiting and Raftery, 2007; Hersbach, 2000). Its analog for multivariate quantities is the energy score (ES) (Gneiting et al., 2008), which in case of an ensemble forecast is computed via

$$\text{ES}(\mathbf{x}_1, \dots, \mathbf{x}_N; \mathbf{y}) = \frac{1}{N} \sum_{n=1}^N \|\mathbf{x}_n - \mathbf{y}\| - \frac{1}{2N^2} \sum_{\nu=1}^N \sum_{n=1}^N \|\mathbf{x}_\nu - \mathbf{x}_n\|,$$

where $\|\cdot\|$ denotes the Euclidean norm. The ES might fail to detect misspecifications in the correlation structure (Pinson and Tastu, 2013; Scheuerer and Hamill, 2015). In contrast, the variogram score (VS) (Scheuerer and Hamill, 2015) is more sensitive in this respect. Here, we use the VS in the form


$$\text{VS}(\mathbf{x}_1, \dots, \mathbf{x}_N; \mathbf{y}) = \sum_{\ell=1}^L \sum_{\lambda=1}^L \left(|y_\ell - y_\lambda|^{\frac{1}{2}} - \frac{1}{N} \sum_{n=1}^N |x_n^\ell - x_n^\lambda|^{\frac{1}{2}} \right)^2,$$

being the special case of a 0.5-VS for ensemble forecasts in the general formulation in Scheuerer and Hamill (2015), which additionally includes optional non-negative weights. However, as inter-variable dependencies, for which a weighting is not obvious or makes sense (Scheuerer and Hamill, 2015), are part of our considerations, we omit the weights here.


In our case study, we report average scores over all forecast cases within a specific test period. As we deal with weather variables having distinct units and magnitudes (wind speed and temperature), the forecasts and observations are standardized for an evaluation via scores here. Precisely, a quantity z is transformed to a standardized quantity z^* via $z^* := (z - \mu)/\sigma$, where μ is taken to be the climatological mean over all available observations in the data set for the respective quantity at the specific location, and σ the corresponding climatological standard deviation.

4.3 Results

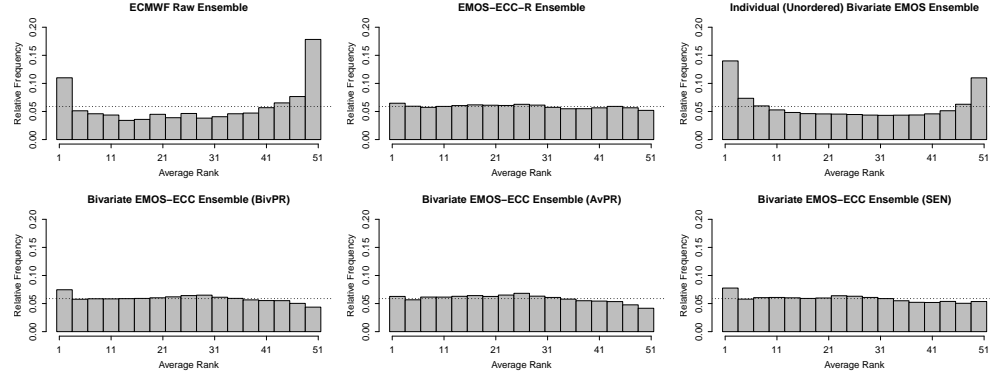
The average energy scores (ES) and variogram scores (VS), respectively, as overall performance measures for the different ensembles are shown in Table 1. Calibration is checked via the multivariate, band depth and average rank histograms, respectively, in Figure 3, with the corresponding reliability indices in Table 2. To account for the random component in the sampling procedures, the scores for the EMOS-ECC-R, the Individual (Unordered) Bivariate EMOS and the three Bivariate EMOS-ECC



(a) Multivariate rank histograms



(b) Band depth rank histograms



(c) Average rank histograms

Figure 3: (a) Multivariate, (b) band depth and (c) average rank histograms for 24 hour ahead temperature and wind speed forecasts at Vienna, Bratislava and Budapest jointly over 3983 test days during the period from 1 January 2003 to 31 December 2013. The results for all ensembles except for the raw ensemble are aggregated over 200 different samples.

Table 1: Average energy scores (ES) and variogram scores (VS) for 24 hour ahead temperature and wind speed forecasts at Vienna, Bratislava and Budapest jointly over 3983 test days during the period from 1 January 2003 to 31 December 2013. The scores for all ensembles except for the raw ensemble are averaged over 200 different samples.

	ES	VS
ECMWF Raw Ensemble	1.014	3.348
EMOS-ECC-R Ensemble	0.773	2.467
Individual (Unordered) Bivariate EMOS Ensemble	0.776	2.471
Bivariate EMOS-ECC Ensemble (BivPR)	0.775	2.440
Bivariate EMOS-ECC Ensemble (AvPR)	0.775	2.439
Bivariate EMOS-ECC Ensemble (SEN)	0.775	2.437

ensembles are aggregated and reported as averages, respectively, over 200 different samples. As our focus is on the multivariate evaluation, we do not explicitly show results for the univariate EMOS postprocessing, but mention briefly that EMOS clearly outperforms the raw ensemble, both for temperature and wind speed.

With respect to the multivariate evaluation, the raw ensemble exhibits a U-shaped multivariate and average rank histogram and a left-skewed band depth rank histogram, each with a strong deviation from uniformity. It is thus clearly underdispersive and consequently uncalibrated, hence calling for postprocessing. Indeed, all postprocessed ensembles outperform the raw ensemble in terms of scores and reliability indices.

The four postprocessed ensembles respecting the dependence pattern of the raw ensemble (that is, the EMOS-ECC-R and the three Bivariate EMOS-ECC ensembles) perform well. In contrast, the postprocessed Individual (Unordered) Bivariate EMOS ensemble not accounting for dependencies reveals weaknesses, although it also outperforms the raw ensemble. Considering the ES only, these shortcomings do not become that visible, as the differences between the ES values of the postprocessed ensembles are rather minor. While the ES is suitable to emphasize the benefits of postprocessing itself, it fails to detect misspecifications in the correlation structures (Pinson and Tastu, 2013; Scheuerer and Hamill, 2015). The VS has shown to be more sensitive in this regard (Scheuerer and Hamill, 2015). Indeed, in terms of the VS, all postprocessed ensembles respecting dependencies perform better than the Individual (Unordered) Bivariate EMOS ensemble. The unsatisfactory performance of the Individual (Unordered) Bivariate EMOS ensemble is further witnessed by the U-shape of the corresponding band depth and average rank histogram, respectively, which indicates an underestimation of the correlation structure (Thorarinsdottir et al., 2015), also resulting in poor reliability indices corresponding to these histograms. Note that with respect to the multivariate rank histogram, the Individual (Unordered) Bivariate EMOS ensemble admittedly performs even best, showing the lowest reliability index, if only slightly. However, similar to the ES, the multivariate rank histogram does not appear to be suitable to detect inadequate correlation structures, especially in higher dimensions (Thorarinsdottir et al., 2015). In a nutshell, the ECC-like reordering proves to be indeed essential in our new approach.

Table 2: Reliability indices Δ_{MR} , Δ_{BDR} and Δ_{AvR} corresponding to the multivariate, band depth and average rank histograms, respectively, in Figure 3, evaluating 24 hour ahead temperature and wind speed forecasts at Vienna, Bratislava and Budapest jointly over 3983 test days during the period from 1 January 2003 to 31 December 2013. The results for all ensembles except for the raw ensemble are based on aggregations over 200 different samples.

	Δ_{MR}	Δ_{BDR}	Δ_{AvR}
ECMWF Raw Ensemble	0.622	1.261	0.397
EMOS-ECC-R Ensemble	0.181	0.158	0.052
Individual (Unordered) Bivariate EMOS Ensemble	0.108	0.191	0.306
Bivariate EMOS-ECC Ensemble (BivPR)	0.150	0.121	0.071
Bivariate EMOS-ECC Ensemble (AvPR)	0.138	0.135	0.093
Bivariate EMOS-ECC Ensemble (SEN)	0.111	0.124	0.076

Concerning the three Bivariate EMOS-ECC ensembles in comparison among themselves, each ensemble performs well in principle, and there are no differences in terms of the ES. With respect to the VS, the three Bivariate EMOS-ECC ensemble perform well on nearly the same level, with the Bivariate EMOS-ECC (SEN) ensemble being slightly best. These findings are also reflected in the rank histograms and the discrepancy indices. As far as band depth and average rank histograms are concerned, the three Bivariate EMOS-ECC ensembles perform more or less equally well, while in the case of the multivariate rank histogram, the Bivariate EMOS-ECC (SEN) ensemble performs better than the other two Bivariate EMOS-ECC ensembles in a more pronounced way. This appears to be plausible, as the ranking according to the signed Euclidean norm has been specifically tailored to the situation in this case study, whereas the rankings following the multivariate or average pre-rank notions are more general, but as shown nevertheless appropriate.

With respect to the ES, the EMOS-ECC-R ensemble relying on univariate EMOS postprocessing very slightly outperforms the three Bivariate EMOS-ECC ensembles based on bivariate EMOS postprocessing. In contrast, the Bivariate EMOS-ECC ensembles outperform EMOS-ECC-R in a more pronounced way in terms of the VS being more sensitive to correlations. Furthermore, the Bivariate EMOS-ECC ensembles exhibit a bit more skill than the EMOS-ECC-R ensemble considering the multivariate and band depth rank histograms and the corresponding reliability indices. In our case study here, the benefits of using bivariate EMOS instead of univariate EMOS postprocessing are apparent, but may not come into effect that much, as the correlations between wind speed and temperature are existent, but not severely pronounced at each of the single locations. In contrast, the spatial correlations between Vienna, Bratislava and Budapest are much stronger for both weather variables, such that the ECC-like reordering plays a more important role here than the bivariate modeling. Moreover, recall that the bivariate EMOS postprocessing has been performed by optimizing with respect to the logarithmic score, while the univariate EMOS methods have employed an optimization in terms of the CRPS, which usually yields better scores. If there was an efficient way to optimize with respect to the CRPS in the case of a bivariate truncated normal distribution, for which no closed-form expression for the CRPS

is available, the results for the bivariate EMOS-based approaches could likely be improved further. In summary, the bivariate EMOS modeling makes sense and is useful here, but the need therefore would be again more important if the inter-variable correlations were more pronounced. Note also that the bivariate EMOS method does not require considerably more computational effort than the univariate EMOS postprocessing for our purposes here.

5 Discussion

We have shown how to combine parametric multi-(low-)dimensional ensemble postprocessing approaches with non-parametric reordering notions based on dependence templates, leading to the novel LDP-Reordering method.

In this context, the challenge of defining a multivariate ranking arises, and we have considered three options to address this issue: the multivariate pre-rank (MultPR), the average pre-rank (AvPR) and the signed Euclidean norm (SEN). The MultPR- and AvPR-based rankings are rather general and may often be subject to a random component due to the random allocation of possible ties. In contrast, the SEN-based ranking concept has been specifically designed for the frame of wind speed and temperature forecasts as examined in our case study and is likely to involve fewer ties than the other two concepts and thus confronted with less randomization. More sophisticated ranking concepts for multivariate settings, maybe tailored to specific settings one is interested in, and/or more sophisticated allocation methods in case of ties could be developed in the future.

Moreover, further research about how to sample from a multivariate distribution is strongly encouraged. In the univariate case, equidistant quantiles provide a canonical and to some extent also optimal choice when drawing a sample (Bröcker, 2012; Graf and Luschgy, 2000). However, similar optimality results are not available for the multivariate situation. As a first attempt into this direction, one could for instance seek a generalization of the concept of equidistance for multidimensional settings, which is however not obvious at first sight, and check for optimality.

From a mathematical point of view, it remains to explicitly relate the LDP-Reordering approach to the (empirical) copula framework, similarly as shown for the univariate reordering approaches (Scheffzik, 2015a). In the LDP-Reordering method, we seek a multivariate predictive CDF with fixed multivariate margins. If all marginal CDFs F_1, \dots, F_L are univariate, the celebrated Sklar's theorem (Nelsen, 2006; Sklar, 1959) allows for a construction of a multivariate CDF H using a suitable copula function C , which models the dependence, via the decomposition

$$H(y_1, \dots, y_L) = C(F_1(y_1), \dots, F_L(y_L))$$

for $y_1, \dots, y_L \in \mathbb{R}$. It would be nice if a result similar to Sklar's theorem could be stated in the case of multivariate (not necessarily univariate) margins, but this appears to be a non-trivial open question for further research. For some special cases, such as for fixed bi- or trivariate margins, theory is available to some extent (Joe, 1997, 2014).

From an applied perspective, we have employed the bivariate EMOS approach of Baran and Möller (2015b) to postprocess wind speed and temperature ensemble forecasts jointly, but at individual stations. The spatial correlations have been addressed by a reordering in an ECC-like manner using the raw ensemble forecasts as a dependence template, based on the three different multivariate ranking

concepts mentioned above. All Bivariate EMOS-ECC ensembles have shown a good predictive performance. The benefits of using both a multivariate reordering and a bi- or multivariate parametric postprocessing to obtain the marginal distributions have become clear in our case study. However, the latter step may be even more powerful in other settings in which the inter-variable correlation structures are more pronounced. For instance, one could work with the bivariate EMOS approach of Schuhen et al. (2012) to postprocess (u, v) -wind forecast vectors or with the Gaussian copula method of Möller et al. (2013) to postprocess temperature and pressure forecasts jointly, as for both pairs of variables, the dependence patterns might be expected to be stronger. The design of further case studies to examine this on the basis of a suitable data set is highly desirable. Another further step is to perform case studies using an observation-based multivariate reordering according to the Schaake shuffle (Clark et al., 2004) or modifications thereof (Scheffzik, 2015b), instead of employing the ECC-like reordering based on a template of raw forecasts.

Acknowledgments

I gratefully acknowledge the support of the Klaus Tschira Foundation and the support of the VolkswagenStiftung under the project “Mesoscale Weather Extremes: Theory, Spatial Modeling and Prediction”. Initial work on this paper was done during my time as Ph.D. student at Heidelberg University, funded by Deutsche Forschungsgemeinschaft through Research Training Group (RTG) 1953. I thank Tilmann Gneiting for valuable comments and suggestions. The forecast data used in the case study have been made available by the European Centre for Medium-Range Weather Forecasts. I thank Stephan Hemri for help with the data, and Sándor Baran and Michael Scheuerer for R code.

References

- J. L. Anderson. A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Clim.*, 9:1518–1530, 1996. doi: 10.1175/1520-0442(1996)009<1518:AMFPAE>2.0.CO;2.
- S. Baran. Probabilistic wind speed forecasting using Bayesian model averaging with truncated normal components. *Comput. Stat. Data Anal.*, 75:227–238, 2014. doi: 10.1016/j.csda.2014.02.013.
- S. Baran and S. Lerch. Log-normal distribution based EMOS models for probabilistic wind speed forecasting. *Q. J. R. Meteorol. Soc.*, 141:2289–2299, 2015. doi: 10.1002/qj.2521.
- S. Baran and A. Möller. Joint probabilistic forecasting of wind speed and temperature using Bayesian model averaging. *Environmetrics*, 26:120–132, 2015a. doi: 10.1002/env.2316.
- S. Baran and A. Möller. Bivariate ensemble model output statistics approach for joint forecasting of wind speed and temperature. *arXiv:1507.03479*, 2015b. Accessed 17 December 2015.
- V. J. Berrocal, A. E. Raftery, and T. Gneiting. Combining spatial statistical and ensemble information in probabilistic weather forecasts. *Mon. Weather Rev.*, 135:1386–1402, 2007. doi: 10.1175/MWR3341.1.

- V. J. Berrocal, A. E. Raftery, and T. Gneiting. Probabilistic quantitative precipitation field forecasting using a two-stage spatial model. *Ann. Appl. Stat.*, 2:1170–1193, 2008. doi: 10.1214/08-AOAS203.
- J. Bröcker. Evaluating raw ensembles with the continuous ranked probability score. *Q. J. R. Meteorol. Soc.*, 138:1611–1617, 2012. doi: 10.1002/qj.1891.
- R. Buizza. The ECMWF ensemble prediction system. In T. N. Palmer and R. Hagedorn, editors, *Predictability of Weather and Climate*, pages 459–489. Cambridge University Press, Cambridge, 2006.
- M. P. Clark, S. Gangopadhyay, L. E. Hay, B. Rajagopalan, and R. L. Wilby. The Schaake shuffle: A method for reconstructing space-time variability in forecasted precipitation and temperature fields. *J. Hydrometeorol.*, 5:243–262, 2004. doi: 10.1175/1525-7541(2004)005<0243:TSSAMF>2.0.CO;2.
- L. Delle Monache, J. P. Hacker, Y. Zhou, X. Deng, and R. B. Stull. Probabilistic aspects of meteorological and ozone regional ensemble forecasts. *J. Geophys. Res.*, 111:D24307, 2006. doi: 10.1029/2005JD006917.
- ECMWF Directorate. Describing ECMWF’s forecasts and forecasting system. *ECMWF Newslett.*, 133:11–13, 2012.
- K. Feldmann, M. Scheuerer, and T. L. Thorarinsdottir. Spatial postprocessing of ensemble forecasts for temperature using nonhomogeneous Gaussian regression. *Mon. Weather Rev.*, 143:955–971, 2015. doi: 10.1175/MWR-D-14-00210.1.
- T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310:248–249, 2005. doi: 10.1126/science.1115255.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.*, 102:359–378, 2007. doi: 10.1198/016214506000001437.
- T. Gneiting, A. E. Raftery, A. H. Westveld, and T. Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Weather Rev.*, 133:1098–1118, 2005. doi: 10.1175/MWR2904.1.
- T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. B*, 69:243–268, 2007. doi: 10.1111/j.1467-9868.2007.00587.x.
- T. Gneiting, L. I. Stanberry, E. P. Gritti, L. Held, and N. A. Johnson. Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test*, 17:211–235, 2008. doi: 10.1007/s11749-008-0114-x.
- D. Gombos, J. A. Hansen, J. Du, and J. McQueen. Theory and applications of the minimum spanning tree rank histogram. *Mon. Weather Rev.*, 135:1490–1505, 2007. doi: 10.1175/MWR3362.1.
- I. J. Good. Rational decisions. *J. R. Stat. Soc. B*, 14:107–114, 1952.
- S. Graf and H. Luschgy. *Foundations of Quantization for Probability Distributions*. Lecture Notes in Mathematics 1730, Springer, Berlin, 2000.

- T. M. Hamill. Interpretation of rank histograms for verifying ensemble forecasts. *Mon. Weather Rev.*, 129:550–560, 2001. doi: 10.1175/1520-0493(2001)129<0550:IORHFV>2.0.CO;2.
- T. M. Hamill and S. J. Colucci. Verification of Eta-RSM short-range ensemble forecasts. *Mon. Weather Rev.*, 125:1312–1327, 1997. doi: 10.1175/1520-0493(1997)125<1312:VOERSR>2.0.CO;2.
- H. Hersbach. Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, 15:559–570, 2000. doi: 10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2.
- H. Joe. *Multivariate Models and Dependence Concepts*. Chapman and Hall, London, 1997.
- H. Joe. *Dependence Modeling with Copulas*. CRC Press, Boca Raton, 2014.
- S. Lerch and T. L. Thorarinsdottir. Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. *Tellus A*, 65:21206, 2013. doi: 10.3402/tellusa.v65i0.21206.
- M. Leutbecher and T. N. Palmer. Ensemble forecasting. *J. Comput. Phys.*, 227:3515–3539, 2008. doi: 10.1016/j.jcp.2007.02.014.
- A. Möller, A. Lenkoski, and T. L. Thorarinsdottir. Multivariate probabilistic forecasting using ensemble Bayesian model averaging and copulas. *Q. J. R. Meteorol. Soc.*, 139:982–991, 2013. doi: 10.1002/qj.2009.
- F. Molteni, R. Buizza, T. N. Palmer, and T. Petroliaxis. The new ECMWF ensemble prediction system: Methodology and validation. *Q. J. R. Meteorol. Soc.*, 122:73–119, 1996. doi: 10.1002/qj.49712252905.
- R. B. Nelsen. *An Introduction to Copulas*. Springer, New York, 2nd edition, 2006.
- T. N. Palmer. The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Q. J. R. Meteorol. Soc.*, 128:747–774, 2002. doi: 10.1256/0035900021643593.
- P. Pinson. Adaptive calibration of (u, v) -wind ensemble forecasts. *Q. J. R. Meteorol. Soc.*, 138:1273–1284, 2012. doi: 10.1002/qj.1873.
- P. Pinson and J. Tastu. Discrimination ability of the Energy score. Technical report, Technical University of Denmark, 2013.
- P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, and B. Klöckl. From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy*, 12:51–62, 2009. doi: 10.1002/we.284.
- R Core Team. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL <http://www.R-project.org>. Accessed 17 December 2015.
- A. E. Raftery, T. Gneiting, F. Balabdaoui, and M. Polakowski. Using Bayesian model averaging to calibrate forecast ensembles. *Mon. Weather Rev.*, 133:1155–1174, 2005. doi: 10.1175/MWR2906.1.
- M. S. Roulston and L. A. Smith. Evaluating probabilistic forecasts using information theory. *Mon. Weather Rev.*, 130:1653–1660, 2002. doi: 10.1175/1520-0493(2002)130<1653:EPFUIT>2.0.CO;2.

- R. Schefzik. Multivariate discrete copulas, with applications in probabilistic weather forecasting. *Publ. Inst. Stat. Univ. Paris*, 59, fasc. 1-2:87–116, 2015a.
- R. Schefzik. A similarity-based implementation of the Schaake shuffle. [arXiv:1507.02079](https://arxiv.org/abs/1507.02079), 2015b. Accessed 17 December 2015.
- R. Schefzik, T. L. Thorarinsdottir, and T. Gneiting. Uncertainty quantification in complex simulation models using ensemble copula coupling. *Stat. Sci.*, 28:616–640, 2013. doi: 10.1214/13-STS443.
- M. Scheuerer. Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Q. J. R. Meteorol. Soc.*, 140:1086–1096, 2014. doi: 10.1002/qj.2183.
- M. Scheuerer and T. M. Hamill. Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Mon. Weather Rev.*, 143:1321–1334, 2015. doi: 10.1175/MWR-D-14-00269.1.
- C. Schoelzel and A. Hense. Probabilistic assessment of regional climate change in Southwest Germany by ensemble dressing. *Clim. Dyn.*, 36:2003–2014, 2011. doi: 10.1007/s00382-010-0815-1.
- N. Schuhen, T. L. Thorarinsdottir, and T. Gneiting. Ensemble model output statistics for wind vectors. *Mon. Weather Rev.*, 140:3204–3219, 2012. doi: 10.1175/MWR-D-00028.1.
- A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris*, 8: 229–231, 1959.
- J. M. Sloughter, T. Gneiting, and A. E. Raftery. Probabilistic wind vector forecasting using ensembles and Bayesian model averaging. *Mon. Weather Rev.*, 141:2107–2119, 2013. doi: 10.1175/MWR-D-12-00002.1.
- L. A. Smith and J. A. Hansen. Extending the limits of ensemble forecast verification with the minimum spanning tree histogram. *Mon. Weather Rev.*, 132:1522–1528, 2004. doi: 10.1175/1520-0493(2004)132<1522:ETLOEF>2.0.CO;2.
- O. Talagrand, R. Vautard, and B. Strauss. Evaluation of probabilistic prediction systems. In *Proceeding of ECMWF Workshop on Predictability*, pages 1–25. Reading, UK, European Centre for Medium-Range Weather Forecasts, 1997.
- T. L. Thorarinsdottir and T. Gneiting. Probabilistic forecasts of wind speed: Ensemble model output statistics using heteroskedastic censored regression. *J. R. Stat. Soc. A*, 173:371–388, 2010. doi: 10.1111/j.1467-985X.2009.00616.x.
- T. L. Thorarinsdottir, M. Scheuerer, and C. Heinz. Assessing the calibration of high-dimensional ensemble forecasts using rank histograms. *J. Comput. Graph. Stat.*, 2015. To appear.
- S. Wilhelm and B. G. Manjunath. `tmvtnorm`: A package for the truncated multivariate normal distribution. *The R Journal*, 2:25–29, 2010.
- S. Wilhelm and B. G. Manjunath. `tmvtnorm`: Truncated Multivariate Normal and Student t Distribution. R package version 1.4-9, 2014. URL <http://CRAN.R-project.org/package=tmvtnorm>. Accessed 17 December 2015.

- D. S. Wilks. The minimum spanning tree histogram as a verification tool for multidimensional ensemble forecasts. *Mon. Weather Rev.*, 132:1329–1340, 2004. doi: 10.1175/1520-0493(2004)132<1329:TMSTHA>2.0.CO;2.
- D. S. Wilks. *Statistical Methods in the Atmospheric Sciences*. Academic Press, Amsterdam, 3rd edition, 2011.
- D. S. Wilks. Multivariate ensemble Model Output Statistics using empirical copulas. *Q. J. R. Meteorol. Soc.*, 141:945–952, 2015. doi: 10.1002/qj.2414.
- R. A. Yuen, T. Gneiting, T. L. Thorarinsdottir, and C. Fraley. ensembleMOS: Ensemble Model Output Statistics. R package version 0.7, 2013. URL <http://CRAN.R-project.org/package=ensembleMOS>. Accessed 17 December 2015.